Big Data Frameworks: A Comparative Study

Wissem Inoubli
University of Jendouba,
Avenue de l'Union du
Maghreb Arabe
Jendouba 8189, Tunisia

Sabeur Aridhi
Université de Lorraine, LORIA,
UMR 7503,
Vandoeuvre-lès-Nancy,
F-54506, France
INRIA Nancy Grand Est,
Villers-lès-Nancy, F-54600,
France

Haithem Mezni
University of Jendouba,
Avenue de l'Union du
Maghreb Arabe
Jendouba 8189, Tunisia

Alexander Jung
Aalto University, School of
Science
P.O. Box 12200, FI-00076,
Finland

arXiv:1610.09962v1 [cs.DC] 31 Oct 2016

ABSTRACT

Recently, increasingly large amounts of data are generated from a variety of sources. Existing data processing technologies are not suitable to cope with the huge amounts of generated data. Yet, many research works focus on Big Data, a *buzzword* referring to the processing of massive volumes of (unstructured) data. Recently proposed frameworks for Big Data applications help to store, analyze and process the data. In this paper, we discuss the challenges of Big Data and we survey existing Big Data frameworks. We also present an experimental evaluation and a comparative study of the most popular Big Data frameworks. This survey is concluded with a presentation of best practices related to the use of studied frameworks in several application domains such as machine learning, graph processing and real-world applications.

Keywords

Big Data; MapReduce; Hadoop; HDFS; Spark; Flink; Storm

1. INTRODUCTION

In recent decades, increasingly large amounts of data are generated from a variety of sources. The size of generated data per day on the Internet has already exceeded two exabytes [1]. Within one minute, 72 hours of videos are uploaded to Youtube, around 30.000 new posts are created on the Tumblr blog platform, more than 100.000 Tweets are shared on Twitter and more than 200.000 pictures are posted on Facebook [1].

Big Data problems lead to several research questions such as (1) how to design scalable environments, (2) how to provide fault tolerance and (3) how to design efficient solutions. Most existing tools for storage, processing and analysis of data are inadequate for massive volumes of heterogeneous data. Consequently, there is an urgent need for more advanced and adequate Big Data solutions.

Many definitions of Big Data have been proposed throughout the literature. Most of them agreed that Big Data problems share four main characteristics, referred to as the four V's (Volume, Variety, Veracity and Velocity) [2].

- **Volume:** it refers to the size of datasets which typically require distributed storage.
- **Variety:** it refers to the fact that Big Data is composed of several different types of data such as text, sound, image and video.
- **Veracity:** it refers to the biases, noise and abnormality in data.
- **Velocity:** it deals with the pace at which data flows in from various sources like social networks, mobile devices and Internet of Things (IoT).

In this paper, we give an overview of some of the most popular and widely used Big Data frameworks which are designed to cope with the above mentioned Big Data problems. We identify some key features which characterize Big Data frameworks. These key features include the programming model and the capability to allow for iterative processing of (streaming) data. We also give a categorization of existing frameworks according to the presented key features.

More specifically, our contributions are the following:

- We present an overview of most popular Big Data frameworks and we categorize them according to some features.
- We experimentally evaluate the performance of the presented frameworks and we present a comparative study of most popular frameworks.
- We highlight best practices related to the use of popular Big Data frameworks in several application domains.

The remainder of the paper is organized as follows. In Section 2, we present the MapReduce programming model. In Section 3, we discuss existing Big Data frameworks and provide a categorization of them. In Section 4, we present a comparative study of the presented Big Data frameworks and we discuss the obtained results. In Section 5, we present some best practices of the studied frameworks. Some concluding points are given in Section 6.

2. MAPREDUCE PROGRAMMING MODEL

MapReduce is a programming model that was designed to deal with parallel processing of large datasets. MapReduce has been proposed by Google in 2004 [3] as an abstraction that allows to perform simple computations while hiding the details of parallelization, distributed storage, load balancing and enabling fault tolerance. The central features of the MapReduce programming model are two functions, written by a user: Map and Reduce. The Map function takes as input a single key-value pair and produces a list of intermediate key-value pairs. The intermediate values associated with the same intermediate key are grouped together and passed to the Reduce function. The Reduce function takes as input an intermediate key and a set of values for that key. It merges these values together to form a smaller set of values. The system overview of MapReduce is illustrated in Fig. 1.

As shown in Fig. 1, the basic steps of a MapReduce program are as follows:

1. **Data reading:** in this phase, the input data is transformed to a set of key-value pairs. The input data may come from various sources such as file system, database management system (DBMS) or main memory (RAM). The input data is split into several fixed-size chunks. Each chunk is processed by one instance of the Map function.
2. **Map phase:** for each chunk having the key-value structure, the corresponding Map function is triggered and produces a set of intermediate key-value pairs.
3. **Combine phase:** this step aims to group together all intermediate key-value pairs associated with the same intermediate key.
4. **Partitioning phase:** following their combination, the results are distributed across the different Reduce functions.
5. **Reduce phase:** the Reduce function merges key-value pairs having the same key and computes a final result.

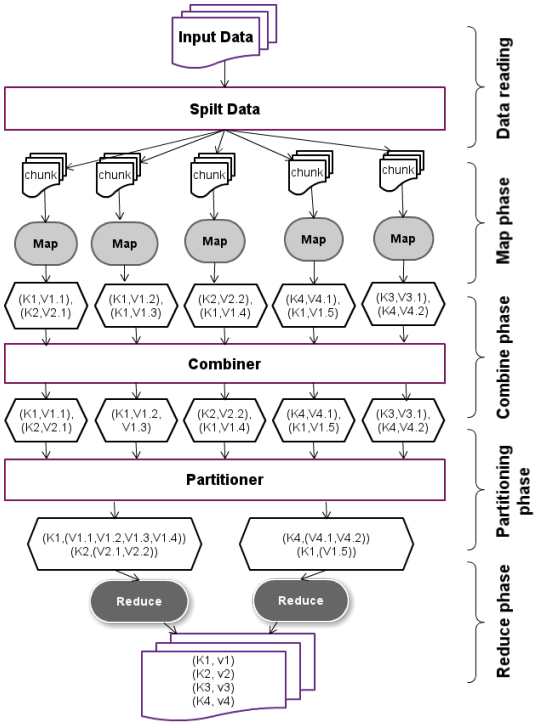

Figure 1: The MapReduce architecture.

Although MapReduce has proven to be effective, it seems that the related programming model is not optimal and generic. By studying the literature, we find that several studies and various open source projects are concentrated around MapReduce to improve its performance. In this context, challenges around this framework can be classified into two main area: (1) data management and (2) scheduling.

Data management. In order to initiate MapReduce jobs, we need to define data structures and execution workflows in which the data will be processed and managed. Formerly, traditional relational database management systems were commonly used. However, such systems have shown their disability to scale with significant amounts of data. To meet this need, a wide variety of different database technologies such as distributed file systems and NoSQL databases have been proposed [4] [5]. Nevertheless, the choice of the adequate system for a specific application remains challenging [4].

Scheduling. Scheduling is an important aspect of MapReduce [6]. During the execution of a MapReduce job, several scheduling decisions need to be taken. These decisions need to consider several information such as data location, resource availablity and others. For example, during the execution of a MapReduce job, the scheduler tries to overcome problems such as *Map-Skew* which refers to imbalanced computational load among map tasks [7].

3. BIG DATA FRAMEWORKS

In this section, we survey some popular Big Data frameworks and categorize them according to their key features. These

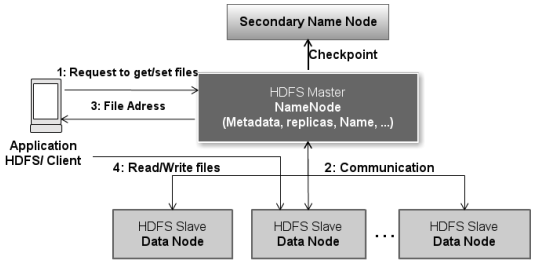

Figure 2: HDFS architecture.

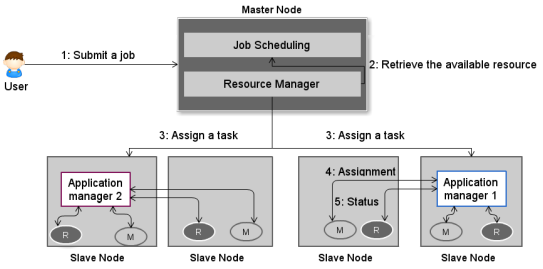

Figure 3: YARN architecture.

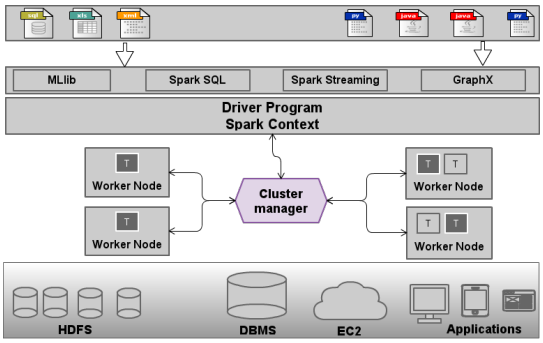

Figure 4: SPARK system overview.

key features are (1) the programming model, (2) the supported programming languages, (3) the type of data sources and (4) the capability to allow for iterative data processing, thereby allowing to cope with streaming data.

3.1 Apache Hadoop

HADOOP is an Apache project founded in 2008 by Doug Cutting at Yahoo and Mike Cafarella at the University of Michigan [8]. HADOOP consists of two main components: (1) Hadoop Distributed File System (HDFS) for data storage and (2) HADOOP MAPREDUCE, an implementation of the MAPREDUCE programming model [3]. In what follows, we discuss HDFS and HADOOP MAPREDUCE.

HDFS. HDFS is an open source implementation of the distributed Google File System (GFS) [5]. It provides a scalable distributed file system for storing large files over distributed machines in a reliable and efficient way [9]. In Fig. 2, we show the abstract architecture of HDFS and its components. It consists of a master/slave architecture with a Name Node being master and several Data Nodes as slaves. The Name Node is responsible for allocating physical space to store large files sent by the HDFS client. If the client wants to retrieve data from HDFS, it sends a request to the Name Node. The Name Node will seek their location in its indexing system and subsequently sends their address back to the client. The Name Node returns to the HDFS client the meta data (filename, file location, etc.) related to the stored files. A secondary Name Node periodically saves the state of the Name Node. If the Name Node fails, the secondary Name Node takes over automatically.

Hadoop MapReduce. There are two main versions of HADOOP MAPREDUCE. In the first version called MRv1, HADOOP MAPREDUCE is essentially based on two components: (1) the Task Tracker that aims to supervise the execution of the Map/Reduce functions and (2) the Job Tracker which represents the master part and allows resource management and job scheduling/monitoring. The Job Tracker supervises and manages the Task Trackers [9]. In the second version of HADOOP called YARN, the two major functionalities of the Job Tracker have been split into separate daemons: (1) a global Resource Manager and (2) per-application Application Master. In Fig. 3, we illustrate the overall architecture of YARN.

As shown in Fig. 3, the Resource Manager receives and runs MAPREDUCE jobs. The per-application Application Master obtains resources from the ResourceManager and works with the Node Manager(s) to execute and monitor the tasks. In YARN, the Resource Manager (respectively the Node Manager) replaces the Job Tracker (respectively the Task Tracker) [10].

3.2 Apache Spark

Apache SPARK is a powerful processing framework that provides an ease of use tool for efficient analytics of heterogeneous data. It was originally developed at UC Berkeley in 2009 [11]. SPARK has several advantages compared to other Big Data frameworks like HADOOP and STORM. SPARK is used by many companies such as Yahoo, Baidu, and Tencent. A key concept of SPARK is Resilient Distributed Datasets (RDDs). An RDD is basically an immutable collection of objects spread across a SPARK cluster. In SPARK, there are two types of operations on RDDs: (1) transformations and (2) actions. Transformations consist in the creation of new RDDs from existing ones using functions like map, filter, union and join. Actions consist of final result of RDD computations.

In Fig. 4, we present an overview of the SPARK architecture. A SPARK cluster is based on a master/slave architecture with three main components:

- **Driver Program:** this component represents the slave node in a SPARK cluster. It maintains an object called SparkContext that manages and supervises running applications.
- **Cluster Manager:** this component is responsible for orchestrating the workflow of application assigned by Driver Program to workers. It also controls and supervises all resources in the cluster and returns their state to the Driver Program.

- **Worker Nodes:** each Worker Node represents a container of one operation during the execution of a SPARK program.

SPARK offers several Application Programming Interfaces (APIs) [11]:

- **Spark Core:** SPARK Core is the underlying general execution engine for the SPARK platform. All other functionalities and extensions are built on top of it. SPARK Core provides in-memory computing capabilities and a generalized execution model to support a wide variety of applications, as well as Java, Scala, and Python APIs for ease of development.
- **Spark Streaming:** SPARK Streaming enables powerful interactive and analytical applications across both streaming and historical data, while inheriting SPARK's ease of use and fault tolerance characteristics. It can be used with a wide variety of popular data sources including HDFS, Flume [12], KAFKA [13], and Twitter [11].
- **Spark SQL:** SPARK offers a range of features to structure data retrieved from several sources. It allows subsequently to manipulate them using the SQL language [14].
- **Spark MLLib:** SPARK provides a scalable machine learning library that delivers both high-quality algorithms (e.g., multiple iterations to increase accuracy) and high speed (up to 100x faster than MAPREDUCE) [11].
- **GraphX:** GRAPHX [15] is a SPARK API for graph-parallel computation (e.g., PageRank algorithm and collaborative filtering). At a high-level, GRAPHX extends the SPARK RDD abstraction by introducing the Resilient Distributed Property Graph: a directed multigraph with properties attached to each vertex and edge. To support graph computation, GRAPHX provides a set of fundamental operators (e.g., subgraph, joinVertices, and mapReduceTriplets) as well as an optimized variant of the Pregel API [16]. In addition, GRAPHX includes a growing collection of graph algorithms (e.g., PageRank, Connected components, Label propagation and Triangle count) to simplify graph analytics tasks.

3.3 Apache Storm

STORM [17] is an open source framework for processing large structured and unstructured data in real time. STORM is a fault tolerant framework that is suitable for real time data analysis, machine learning, sequential and iterative computation. Following a comparative study of STORM and HADOOP, we find that the first is geared for real time applications while the second is effective for batch applications.

As shown in Fig. 5, a STORM program/topology is represented by a directed acyclic graphs (DAG). The edges of the program DAG represent data transfer. The nodes of the DAG are divided into two types: spouts and bolts. The spouts (or entry points) of a STORM program represent the data sources. The bolts represent the functions to be performed on the data. Note that STORM distributes bolts across multiple nodes to process the data in parallel.

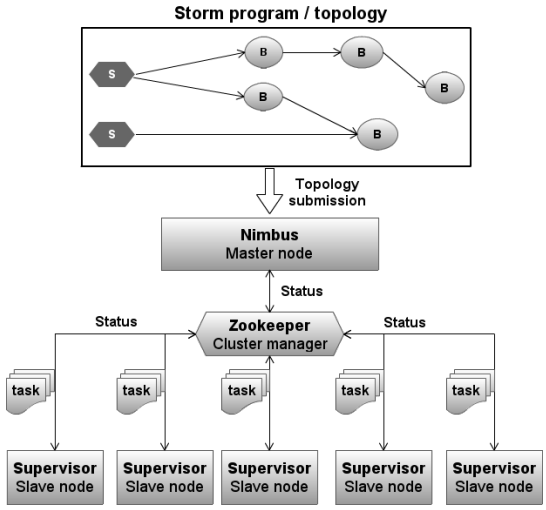

Figure 5: Topology of a STORM program and STORM architecture.

In Fig. 5, we show a STORM cluster administrated by ZOOKEEPER, a service for coordinating processes of distributed applications [18]. STORM is based on two daemons called Nimbus (in master node) and supervisor (for each slave node). Nimbus supervises the slave nodes and assigns tasks to them. If it detects a node failure in the cluster, it re-assigns the task to another node. Each supervisor controls the execution of its tasks (affected by the nimbus). It can stop or start the spots following the instructions of Nimbus. Each topology submitted to STORM cluster is divided into several tasks.

3.4 Apache Flink

FLINK [19] is an open source framework for processing data in both real time mode and batch mode. It provides several benefits such as fault-tolerant and large scale computation. The programming model of FLINK is similar to MAPREDUCE. By contrast to MAPREDUCE, FLINK offers additional high level functions such as join, filter and aggregation. FLINK allows iterative processing and real time computation on stream data collected by different tools such as Flume [12] and KAFKA [13]. It offers several APIs on a more abstract level allowing the user to launch distributed computation in a transparent and easy way. FLINK ML is a machine learning library that provides a wide range of learning algorithms to create fast and scalable Big Data applications. In Fig. 6, we illustrate the architecture and components of FLINK.

As shown in Fig. 6, the FLINK system consists of several layers. In the highest layer, users can submit their programs written in Java or Scala. User programs are then converted by the FLINK compiler to DAGs. Each submitted job is represented by a graph. Nodes of the graph represent operations (e.g., map, reduce, join or filter) that will be applied to process the data. Edges of the graph represent the flow of data between the operations. A DAG produced

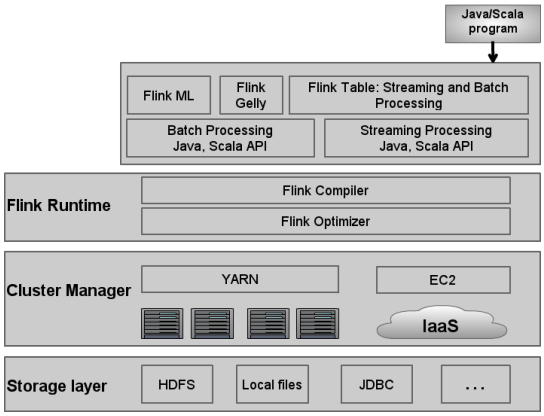

Figure 6: FLINK architecture

by the FLINK compiler is received by the FLINK optimizer in order to improve performance by optimizing the DAG (e.g., re-ordering of the operations). The second layer of FLINK is the cluster manager which is responsible for planning tasks, monitoring the status of jobs and resource management. The lowest layer is the storage layer that ensures storage of the data to multiple destinations such as HDFS and local files.

3.5 Categorization of Big Data Frameworks

We present in Table 1 a categorization of the presented frameworks according to data format, processing mode, used data sources, programming model, supported programming languages, cluster manager and whether the framework allows iterative computation or not.

We mention that HADOOP is currently one of the most widely used parallel processing solutions. HADOOP ecosystem consists of a set of tools such as FLUME, HBASE, HIVE and MAHOUT. HADOOP is widely adopted in the management of large-size clusters. Its YARN daemon makes it a suitable choice to configure Big Data solutions on several nodes [20]. For instance, HADOOP is used by Yahoo to manage 24 thousands of nodes. Moreover, HADOOP MAPREDUCE is proven to be the best choice to deal with text processing tasks [21]. We also notice that HADOOP stores its data in HDFS and does not allow iterative computation, whereas SPARK, FLINK and STORM allow different data sources and iterative computation.

As shown in Table 1, SPARK importance lies in its in-memory features and micro-batch processing capabilities, especially in iterative and incremental processing [22]. In addition, SPARK offers an interactive tool called SPARK-SHELL which allows to exploit the SPARK cluster in real time. Once interactive applications were created, they may subsequently be executed interactively in the cluster. Although SPARK is known to be the fastest framework due to the concept of RDD, it is not a suitable choice in the case of complex data processing. For this purpose, FLINK may be a good candidate. Its mechanism transforms each task to an optimized graph, which allows it to surpass SPARK in the case of complex jobs.

FLINK shares similarities and characteristics with SPARK. It

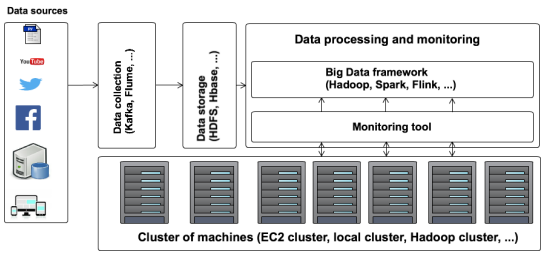
Data sources
Data collection (Kafka, Flume, ...)
Data storage (HDFS, Hbase, ...)
Data processing and monitoring
Big Data framework (Hadoop, Spark, Flink, ...)
Monitoring tool
Cluster of machines (EC2 cluster, local cluster, Hadoop cluster, ...)

Figure 7: *Batch Mode* scenario

offers good processing performance when dealing with complex Big Data structures such as graphs. Although there exist other solutions for large-scale graph processing, FLINK and SPARK are enriched with specific APIs and tools for machine learning, predictive analysis and graph stream analysis [19] [11].

From a technical point of view, we mention that all the presented frameworks support several programming languages like Java, Scala and Python.

4. EXPERIMENTS

We have performed an extensive set of experiments to highlight the strengths and weaknesses of popular Big Data frameworks. The performed analysis covers the performance, scalability, and the resource usage. For our experiments, we evaluated SPARK, HADOOP, FLINK and STORM. Implementation details and benchmarks can be found in the following link: https://members.loria.fr/SAridhi/files/bigdata/. In this section, we first present our experimental environment and our experimental protocol. Then, we discuss the obtained results.

4.1 Experimental environment

All the experiments were performed in a cluster of 10 machines, each equipped with Quad-Core CPU, 8GB of main memory and 500 GB of local storage. For our tests, we used HADOOP 2.6.3, FLINK 0.10.3, SPARK 1.6.0 and STORM 0.9. For all the tested frameworks, we used the default values of their corresponding parameters.

4.2 Experimental protocol

We consider two scenarios according to the data processing mode of the evaluated frameworks.

- In the *Batch Mode* scenario, we evaluate HADOOP, SPARK and FLINK while running the *WordCount* example on a big set of tweets. The used tweets were collected by Apache Flume [12] and stored in HDFS. As shown in Fig. 7, the collected data may come from different sources including social networks, local files, log files and sensors. In our case, Twitter is the main source of our collected data. The motivations behind using Apache Flume to collect the processed tweets is its integration facility in the HADOOP ecosystem (especially the HDFS system). Moreover, Apache FLUME allows data collection in a distributed way and offers high data availability and fault tolerance. We collected 10 billions of tweets and we used them to form large

	Hadoop	Spark	Storm	Flink
Data format	Key-value	RDD	Key-value	Key-value
Processing mode	Batch	Batch and stream	Stream	Batch and stream
Data sources	HDFS	HDFS, DBMS and KAFKA	Spoots	Multiple sources
Programming model	Map and Reduce	Transformation and Action	Bolts	TransformAt
Supported programming language	Java	Java, Scala and Python	Java	Java
Cluster manager	YARN	YARN	YARN or ZOOKEEPER	ZOOKEEPER
Shared resources	Disk	Memory	No shared resources	CPU
Comments	Stores large data in HDFS	Gives several APIs to develop interactive applications	Suitable for real-time applications	FLINK is an extension of MAPREDUCE with graph methods
Iterative computation	No	Yes	Yes	Yes

Table 1: A comparative study of popular Big Data frameworks.

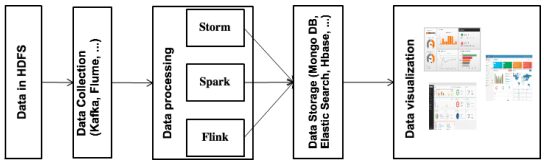

Figure 8: *Stream Mode* scenario

tweet files with a size on disk varying from 250 MB to 40 GB of data.

- In the *Stream Mode* scenario, we evaluate real-time data processing capabilities of STORM, FLINK and SPARK. The *Stream Mode* scenario is divided into three main steps. As shown in Fig. 8, the first step is devoted to data storage. To do this step, we collected 1 billion tweets from Twitter using FLUME and we stored them in HDFS. The stored data is then sent to KAFKA, a messaging server that guarantees fault tolerance during the streaming and message persistence [13]. The second step is responsible for the streaming of tweets to the studied frameworks. To allow simultaneous streaming of the data collected from HDFS by STORM, SPARK and FLINK, we have implemented a script that accesses the HDFS and transfers the data to KAFKA. The last step allows to process the received messages from KAFKA using an Extract, Transform and Load (ETL) routine. Once the messages (tweets) sent by KAFKA are loaded and transformed, they will be stored using ElasticSearch storage server, and possibly visualized with KIBANA [23]. Note that during the ETL steps, incoming messages are processed in a specific time interval for each framework. Regarding the hardware configuration adopted in the *Stream Mode*, we used one machine for KAFKA and one machine for ZOOKEEPER, which allows coordination between KAFKA and STORM. As for processing task, the remaining machines are devoted to access the data in HDFS and to send it to KAFKA server. To determine the number of events (i.e. tweets) for each framework, a short simulation during 15 minutes has led to the following results: 5358 events for SPARK, 9684 events for FLINK, and 10106 events for STORM.

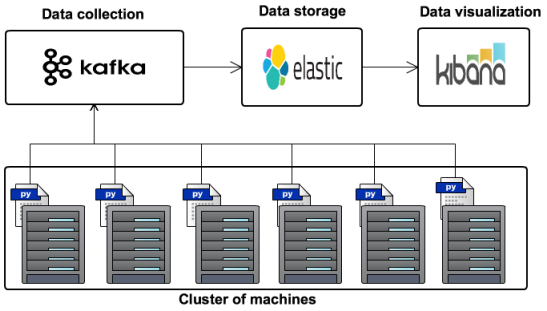

Figure 9: Architecture of our personalized monitoring tool

To monitor resource usage of the studied frameworks, existing tools like AMBARI [24] and Hue [25] are not suitable for our case as they only offer real-time monitoring results. To allow monitoring resources usage according to the executed jobs, we have implemented a personalized monitoring tool as shown in Fig. 9. Our monitoring solution is based on four core components: (1) cluster state detection module, (2) data collection module, (3) data storage module, and (4) data visualization module. To detect the states of the machines, we have implemented a Python script and we deployed it in every machine of the cluster. This script is responsible for collecting CPU, RAM, Disk I/O, and Bandwidth history. The collected data are stored in ElasticSearch in order to be used in the evaluation step. The stored data in ElasticSearch are used by KIBANA for monitoring and visualization purposes.

4.3 Experimental results

4.3.1 Batch mode

In this section, we evaluate the scalability of the studied frameworks. We also measure their CPU, RAM and disk I/O usage and bandwidth consumption while processing tweets, as described in Section 4.2.

Scalability

The first experiment aims to evaluate the impact of data size on the processing time. Experiments are conducted with various datasets with a size on disk varying from 250 MB to 40 GB. Fig. 10 shows the average processing time for each framework and for every dataset. As shown in Fig. 10,

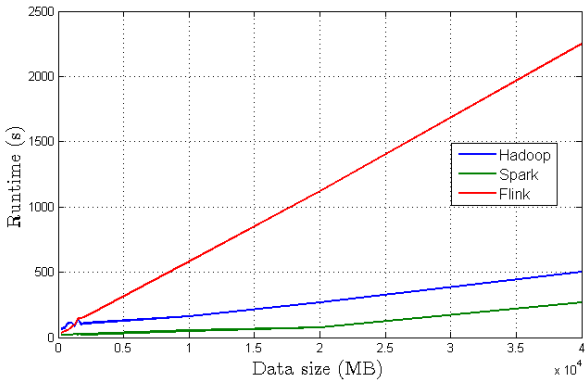

Figure 10: Impact of the data size on the average processing time

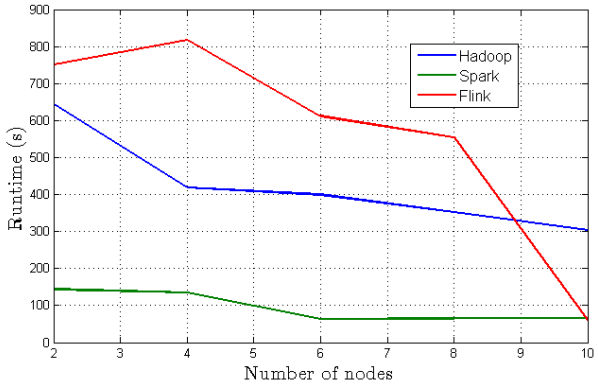

Figure 11: Impact of the number of machines on the average processing time

Spark is the fastest framework for all the datasets, Hadoop is the next and Flink is the lowest. We notice that Flink is faster than Hadoop only in the case of very small datasets (less than 5 GB). Compared to Spark, Hadoop achieves data transfer by accessing the HDFS. Hence, the processing time of Hadoop is considerably affected by the high amount of Input/Output (I/O) operations. By avoiding I/O operations, Spark has gradually reduced the processing time. It can also be observed that the computational time of Flink is longer than the other frameworks. This can be explained by the fact that the access to HDFS increases the time spent to exchange the processed data. Indeed, the outputs of the Mappers are transferred to the Reducers. Thus, extra time is consumed especially for large datasets. We also notice that Spark defines optimal time by reducing the communication between data nodes. However, its behavior is a little mechanical and the processing of large datasets is much slower. This is not the case for Hadoop that can maintain lower evolving in case of larger datasets (e.g. 40 GB in our case).

In the next experiment, we tried to evaluate the scalability and the processing time of the considered frameworks based on the size of the used cluster (the number of the used machines). Fig. 11 gives an overview of the impact of the number of the used machines on the processing time. Both Hadoop and Flink take higher time regardless of the

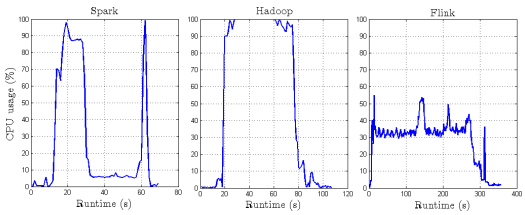

Figure 12: CPU resource usage in *Batch Mode* scenario

cluster size, compared to Spark. Fig. 11 shows instability in the slope of Flink due to the network traffic. In fact, jobs in Flink are modelled as a graph that is distributed on the cluster, where nodes represent Map and Reduce functions, whereas edges denote data flows between Map and Reduce functions. In this case, Flink performance depends on the network state that may affect data/intermediate results from Map to Reduce functions across the cluster. Regarding Hadoop, it is clear that the processing time is proportional to the cluster size. In contrast to reduced number of machines, the gap between Spark and Hadoop is reduced with large cluster. This means that Hadoop performs well and can have close processing time in case of more available resources. The time spent by Spark is approximately between 140 seconds and 145 seconds for 2 to 4 node cluster. Furthermore, as the number of participating nodes increases, the processing time, yet, remains approximately equal to 70 seconds. This is explained by the processing logic of Spark. Indeed, Spark depends on the main memory (RAM) and also the available resources in the cluster. In case of insufficient resources to process the intermediate results, Spark requires more RAM to store those intermediate results. This is the case of 6 to 10 nodes which explain the inability to improve the processing time even with an increased number of participating machines. In the case of Hadoop, intermediate results are stored on disk as there is no storage constraint. This explains the reduced execution time that reached 300 seconds in case of 10 nodes, compared to 410 seconds when exploiting only 4 nodes.

CPU consumption

As shown in Fig. 12, CPU consumption is approximately in direct ratios with the problem scale, especially in the Map phase. However, the slopes of Flink (see Fig. 12) are not larger than those of Spark and Hadoop because Flink partially exploits disk and memory resources unlike Spark and Hadoop. Since Hadoop is initially modelled to frequently use the disk, the amount of Read/Write of the obtained processing results is high. Hence, Hadoop CPU consumption is important. In contrast, Spark mechanism relies on the node memory. This approach is not costly in terms of CPU consumption. Fig. 12 shows Hadoop CPU usage. The processed data are loaded in the first 20 seconds of the total execution time. In the next 60 seconds, Map and Reduce functions are triggered with some gap. When a Reduce function receives the totality of its key-value pairs (intermediate results stored by Map functions on the disk) assigned by the Application Manager, it starts processing immediately. In the remaining processing time (from 80 to 120 seconds), Hadoop writes the final results on disk. As for Spark CPU usage, Spark loads data from disk to the memory during the first 10 seconds. In the next 20 seconds (from 10 to 30

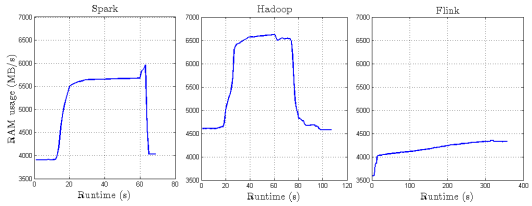

Figure 13: RAM consumption in *Batch Mode* scenario

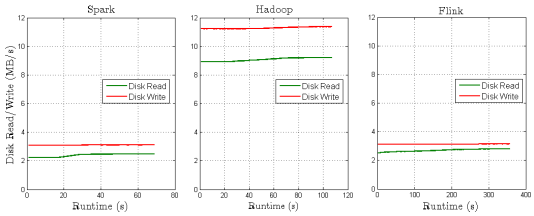

Figure 14: Disk I/O usage in *Batch Mode* scenario

seconds), Map functions are triggered to process the loaded data. In the second half time execution, each Reduce function processes its own data that come from the Map functions. In the last 10 seconds, Spark combines all the results of the Reduce functions. Although Flink execution time is higher than Spark and Hadoop, its CPU consumption is considerably low compared to Spark and Hadoop. Indeed, in the default configuration of Flink, a worker is reserved for each machine, unlike Hadoop which automatically exploits the totality of available resources. This behaviour is seen when several jobs are submitted to the cluster.

RAM consumption

Fig. 13 plots the memory bandwidth of the studied frameworks. The RAM usage rate is almost the same for Hadoop and Spark especially in the Map phase. When data fit into cache, Spark has a bandwidth between 5.62 GB/s and 5.98 GB/s. Spark logic depends mainly on RAM utilization. That is why during 78.57% of the total job execution time (68 seconds), RAM is almost occupied (5.6 GB per node). Regarding Hadoop, only 45.83% of jobs execution time (107 seconds) is used with 6.6 GB of RAM during this time, where 4.6 GB is reserved to the daemons of Hadoop framework. Regarding Flink, RAM is occupied during 97.14% of the total execution time (350 seconds), with 3.6 GB reserved for the daemons which are responsible for managing the cluster. The increase in RAM usage from 4.03 GB to 4.32 GB is explained by the gradual triggering of Reduce functions after the receipt of intermediate results that came from Map functions. Flink models each task as a graph where its constituting nodes represent a specific function and edges denote the flow of data between those nodes. Intermediate results are, hence, sent directly from Map to Reduce functions without exploiting RAM and disk resources.

Disk I/O usage

Fig. 14 shows the amount of disk usage by the three frameworks. Green slopes denote the amount of write, whereas red slopes denote the amount of read. Fig. 14 shows that Hadoop frequently accesses the disk, as the amount of write is about 11 MB/s. This is not the case for Spark (about 3 MB/s) as this framework is mainly memory-oriented. As

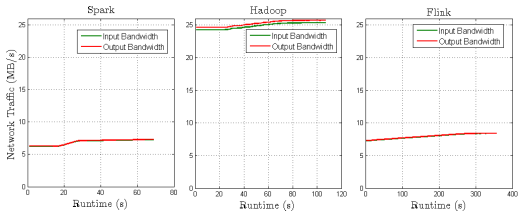

Figure 15: Bandwidth resource usage in *Batch Mode* scenario

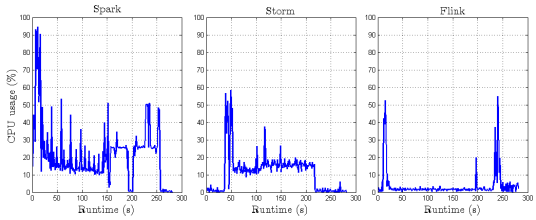

Figure 16: CPU consumption in *Stream Mode* scenario

for Flink which shares a similar behaviour with Spark, disk usage is very low compared to Hadoop (about 3 MB/s). Indeed, Flink uploads, first, the required processing data to the memory and, then, distributes them among the candidate workers.

Bandwidth resource usage

As shown in Fig. 15, Hadoop surpasses Spark and Flink in traffic utilization. The amount of data exchanged per second is high compared to Spark and Flink (7.29 MB/s for Flink, 24.65 MB/s for Hadoop, and 6.26 MB/s for Spark). This is explained by the frequent serialization and migration operations between cluster nodes. Note that Hadoop is based on data locality principle to distribute data between different nodes. Regarding Spark, the processed data are always stored on each worker. In the same fashion, Flink assigns such data to the jobs and candidate workers once in the beginning, which allows reducing data migration between machines.

4.4 Stream Mode

In this experiment, we measure CPU, RAM, disk I/O usage and bandwidth consumption of the studied frameworks while processing tweets, as described in Section 4.2.

CPU consumption

Results have shown that the number of events processed by Storm (10106) is close to that processed by Flink (9684) despite the larger-size nature of events in Flink compared to Storm. In fact, the window time configuration of Storm allows to rapidly deal with the incoming messages. Fig. 16 plots the CPU consumption rate of Flink, Storm and Spark. As shown in Fig. 16, Flink CPU consumption is low compared to Spark and Storm. Flink exploits about 5% of the available CPU to process 9684 events, whereas Storm CPU usage varies between 13% and 16% when processing 10106 events. However, Flink may provide better results than Storm when CPU resources are more exploited. In the literature, Flink is designed to process large messages, unlike Storm which is only able to deal with small messages

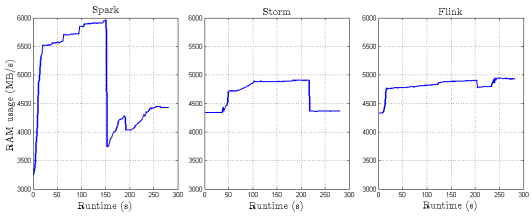

Figure 17: RAM consumption in *Stream Mode* scenario

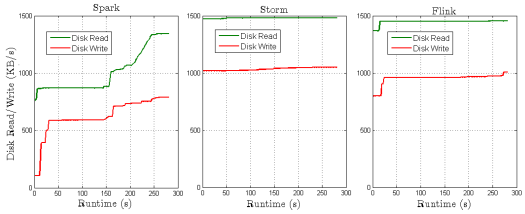

Figure 18: Disk usage in *Stream Mode* scenario

(e.g., messages coming from sensors). This explains the fact that Flink is largely adopted in several applications including smart cities, financial markets, surveillance systems [19]. Unlike Flink and Storm, Spark collects events' data every second and performs processing task after that. Hence, more than one message is processed, which explains the high CPU usage by Spark. Because of Flink's pipeline nature, each message is associated to a thread and consumed at each window time. Consequently, this low volume of processed data does not affect the CPU resource usage.

RAM consumption

Fig. 17 shows the cost of event stream processing in terms of RAM consumption. Spark reached 6 GB (75% of the available resources) due its in-memory behaviour and its ability to perform in micro-batch (process a group of messages at a time). Flink and Storm did not exceed 5 GB (around 62.5% of the available RAM) as their stream mode behaviour consists in processing only single messages. Regarding Spark, the number of processed messages is small. Hence, the communication frequency with the cluster manager is low. In contrast, the number of processed events is high for Flink and Storm, which explains the important communication frequency between the frameworks and their Daemons (i.e. between Storm and Zookeeper, and between Flink and Yarn) is high. Indeed, communication topology in Flink is predefined, whereas the communication topology in the case of Storm is dynamic because NBus (the master component of Storm) searches periodically the available nodes to perform processing tasks.

Disk I/O usage

Fig. 18 depicts the amount of disk usage by the studied frameworks. Red slopes denote the amount of write operations, whereas green slopes denote the amount of read operations. The amounts of write operations in Flink and Storm are almost close. Flink and Storm frequently access the disk and are faster than Spark in terms of the number of processed messages. As discussed in the above sections, Spark framework is an in-memory framework which explains its lower disk usage.

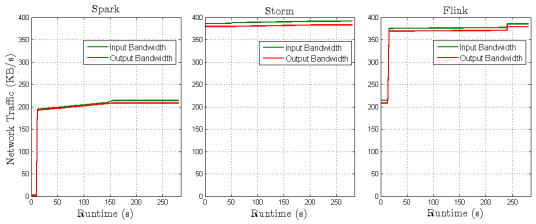

Figure 19: Traffic bandwidth in *Stream Mode* scenario

Bandwidth resource usage

As shown in Fig. 19, the amount of data exchanged per second varies between 375 KB/s and 385 KB/s in case of Flink and varies between 387 KB/s and 390 KB/s in the case of Storm. This amount is high compared to Spark as its bandwidth usage did not exceed 220 KB/s. This is due to the reduced frequency of serialization and migration operations between the cluster nodes as Spark processes a group of messages at each operation. Consequently, the amount of exchanged data is reduced.

4.5 Summary of evaluation

From the runtime experiments it is clear that Spark can deal with complex and large datasets better than Hadoop and Flink. The carried experiments in this work also indicate that Hadoop performs well on the whole. However, it has some limitations regarding the time spent in communicating data between nodes and requires a considerable processing time when the size of data increases. Flink resource consumption is the lowest compared to Spark and Hadoop. This is explained by the greed nature of Spark and Hadoop.

5. BEST PRACTICES

In the previous section, two major processing approaches (batch and stream) were studied and compared in terms of speed and resource usage. Choosing the right processing model is a challenging problem given the growing number of frameworks with similar and various services. This section aims to shed light on the strengths of the above discussed frameworks when exploited in specific fields including stream processing, batch processing, machine learning and graph processing. We also discuss the use of the studied frameworks in several real-world applications including healthcare applications, recommender systems, social network analysis and smart cities.

5.1 Stream processing

As the world becomes more connected and influenced by mobile devices and sensors, stream computing emerged as a basic capability of real-time applications in several domains, including monitoring systems, smart cities, financial markets and manufacturing [22]. However, this flood of data that comes from various sources at high speed always needs to be processed in a short time interval. In this case, Storm and Flink may be considered, as they allow pure stream processing. The design of in-stream applications needs to take into account the frequency and the size of incoming events data. In the case of stream processing, Apache Storm is well-known to be the best choice for the big/high stream oriented applications (billions of events per second/core).

As shown in the conducted experiments, STORM performs well and allows resource saving even if the stream of events becomes important.

5.2 Micro-batch processing

In case of batch processing, Spark may be a suitable framework to deal with periodic processing tasks such as Web usage mining, fraud detection, etc. In some situations, there is a need for a programming model that combines both batch and stream behaviour over the huge volume/frequency of data in a lambda architecture. In this architecture, periodic analysis tasks are performed in a larger window time. Such behaviour is called micro-batch. For instance, data produced by healthcare and IoT applications often require combining batch and stream processing. In this case frameworks like FLINK and SPARK may be good candidates [26]. SPARK micro-batch behaviour allows to process datasets in larger window times. SPARK consists of a set of tools, such as SPARK MLLIB and SPARK Stream that provide rich analysis functionalities in micro-batch. Such behaviour requires regrouping the processed data periodically, before performing analysis task.

5.3 Machine learning algorithms

Machine learning algorithms are iterative in nature [26]. Most of the above discussed frameworks support machine learning capabilities through a set of libraries and APIs. FLINK-ML library includes implementations of k-Means clustering algorithm, logistic regression, and Alternating Least Squares (ALS) for recommendation [27]. SPARK has more efficient set of machine learning algorithms such as SPARK MLLIB and MLI [28]. SPARK MLLIB is a scalable and fast library that is suitable for general needs and most areas of machine learning. Regarding HADOOP framework, Apache Mahout aims to build scalable and performant machine learning applications on top of HADOOP.

5.4 Big Graph processing

The field of processing large graphs has attracted considerable attention because of its huge number of applications, such as the analysis of social networks [29], Web graphs [30] and bioinformatics [31]. It is important to mention that HADOOP is not the optimal programming model for graph processing [32]. This can be explained by the fact that HADOOP uses coarse-grained tasks to do its work, which are too heavyweight for graph processing and iterative algorithms [26]. In addition, HADOOP can not cache intermediate data in memory for faster performance. We also notice that most of Big Data frameworks provide graph-related libraries (e.g., GRAPHX [15] with SPARK and GELLY [33] with FLINK). Moreover, many graph processing systems have been proposed [34]. Such frameworks include PREGEL [16], GRAPHLAB [35], BLADYG [36] and TRINITY [37].

5.5 Real-world applications

5.5.1 Healthcare applications

Healthcare scientific applications, such as body area network provide monitoring capabilities to decide on the health status of a host. This requires deploying hundreds of interconnected sensors over the human body to collect various data including breath, cardiovascular, insulin, blood, glucose and body temperature [38]. However, sending and processing iteratively such stream of health data is not supported by the original MAPREDUCE model. HADOOP was initially designed to process big data already available in the distributed file system. In the literature, many extensions have been applied to the original MAPREDUCE model in order to allow iterative computing such as HALOOP system [39] and Twister [40]. Nevertheless, the two caching functionalities in HALOOP that allow reusing processing data in the later iterations and make checking for a fix-point lack efficiency. Also, since processed data may partially remain unchanged through the different iterations, they have to be reloaded and reprocessed at each iteration. This may lead to resource wastage, especially network bandwidth and processor resources. Unlike HALOOP and existing MAPREDUCE extensions, SPARK provides support for interactive queries and iterative computing. RDD caching makes Spark efficient and performs well in iterative use cases that require multiple treatments on large in-memory datasets [22].

5.5.2 Recommendation systems

Recommender systems is another field that began to attract more attention, especially with the continuous changes and the growing streams of users' ratings [41]. Unlike traditional recommendation approaches that only deal with static item and user data, new emerging recommender systems must adapt to the high volume of item information and the big stream of user ratings and tastes. In this case, recommender systems must be able to process the big stream of data. For instance, news items are characterized by a high degree of change and user interests vary over time which requires a continuous adjustment of the recommender system. In this case, frameworks like HADOOP are not able to deal with the fast stream of data (e.g. user ratings and comments), which may affect the real evaluation of available items (e.g. product or news). In such a situation, the adoption of effective stream processing frameworks is encouraged in order to avoid overrating or incorporating user/item related data into the recommender system. Tools like MAHOUT, FLINK-ML and SPARK MLLIB include collaborative filtering algorithms, that may be used for e-commerce purpose and in some social network services to suggest suitable items to users [42].

5.5.3 Social media

Social media is another representative data source for big data that requires real-time processing and results. Its is generated from a wide range of Internet applications and Web sites including social and business-oriented networks (e.g. LinkedIn, Facebook), online mobile photo and video sharing services (e.g. Instagram, Youtube, Flickr), etc. This huge volume of social data requires a set of methods and algorithms related to, text analysis, information diffusion, information fusion, community detection and network analytics, which may be exploited to analyse and process information from social-based sources [43]. This also requires iterative processing and learning capabilities and necessitates the adoption of in-stream frameworks such as STORM and FLINK along with their rich libraries.

5.5.4 Smart cities

Smart city is a broad concept that encompasses economy, governance, mobility, people, environment and living [44].

It refers to the use of information technology to enhance quality, performance and interactivity of urban services in a city. It also aims to connect several geographically distant cities [45]. Within a smart city, data is collected from sensors installed on utility poles, water lines, buses, trains and traffic lights. The networking of hardware equipment and sensors is referred to as Internet of Things (IoT) and represents a significant source of Big data. Big data technologies are used for several purposes in a smart city including traffic statistics, smart agriculture, healthcare, transport and many others [45]. For example, transporters of the logistic company UPS are equipped with operating sensors and GPS devices reporting the states of their engines and their positions respectively. This data is used to predict failures and track the positions of the vehicles. Urban traffic also provides large quantities of data that come from various sensors (e.g., GPSs, public transportation smart cards, weather conditions devices and traffic cameras). To understand this traffic behaviour, it is important to reveal hidden and valuable information from the big stream/storage of data. Finding the right programming model is still a challenge because of the diversity and the growing number of services [46]. Indeed, some use cases are often slow such as urban planning and traffic control issues. Thus, the adoption of a batch-oriented framework like HADOOP is sufficient. Processing urban data in micro-batch fashion is possible, for example, in case of eGovernment and public administration services. Other use cases like healthcare services (e.g. remote assistance of patients) need decision making and results within few milliseconds. In this case, real-time processing frameworks like STORM are encouraged. Combining the strengths of the above discussed frameworks may also be useful to deal with cross-domain smart ecosystems also called big services [47].

6. CONCLUSIONS

In this work, we surveyed popular frameworks for large-scale data processing. After a brief description of the main paradigms related to Big Data problems, we presented an overview of the Big Data frameworks Hadoop, Spark, Storm and Flink. We presented a categorization of these frameworks according to some main features such as the used programming model, the type of data sources, the supported programming languages and whether the framework allows iterative processing or not. We also conducted an extensive comparative study of the above presented frameworks on a cluster of machines and we highlighted best practices while using the studied Big Data frameworks.

7. REFERENCES

[1] A. Gandomi, M. Haider, Beyond the hype: Big data concepts, methods, and analytics, International Journal of Information Management 35 (2) (2015) 137 – 144.

[2] A. Oguntimilehin, E. Ademola, A review of big data management, benefits and challenges, A Review of Big Data Management, Benefits and Challenges 5 (6) (2014) 433–438.

[3] J. Dean, S. Ghemawat, MapReduce: simplified data processing on large clusters, Commun. ACM 51 (1) (2008) 107–113.

[4] J. R. Lourenço, B. Cabral, P. Carreiro, M. Vieira, J. Bernardino, Choosing the right nosql database for the job: a quality attribute evaluation, Journal of Big Data 2 (1) (2015) 1–26. doi:10.1186/s40537-015-0025-0.

[5] S. Ghemawat, H. Gobioff, S.-T. Leung, The google file system, SIGOPS Oper. Syst. Rev. 37 (5) (2003) 29–43.

[6] R. Vernica, A. Balmin, K. S. Beyer, V. Ercegovac, Adaptive mapreduce using situation-aware mappers, in: Proceedings of the 15th International Conference on Extending Database Technology, EDBT '12, ACM, New York, NY, USA, 2012, pp. 420–431.

[7] Y. Kwon, M. Balazinska, B. Howe, J. Rolia, Skewtune: Mitigating skew in mapreduce applications, in: Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, SIGMOD '12, ACM, New York, NY, USA, 2012, pp. 25–36.

[8] I. Polato, R. Ré, A. Goldman, F. Kon, A comprehensive view of hadoop researchâĂŤa systematic literature review, Journal of Network and Computer Applications 46 (2014) 1–25.

[9] T. White, Hadoop: The definitive guide, " O'Reilly Media, Inc.", 2012.

[10] R. Li, H. Hu, H. Li, Y. Wu, J. Yang, Mapreduce parallel programming model: A state-of-the-art survey, International Journal of Parallel Programming (2015) 1–35.

[11] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, I. Stoica, Spark: Cluster computing with working sets, in: Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing, HotCloud'10, USENIX Association, Berkeley, CA, USA, 2010, pp. 10–10.

[12] C. Chambers, A. Raniwala, F. Perry, S. Adams, R. R. Henry, R. Bradshaw, N. Weizenbaum, Flumejava: Easy, efficient data-parallel pipelines, in: Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '10, ACM, New York, NY, USA, 2010, pp. 363–375.

[13] N. Garg, Apache Kafka, Packt Publishing, 2013.

[14] M. Armbrust, R. S. Xin, C. Lian, Y. Huai, D. Liu, J. K. Bradley, X. Meng, T. Kaftan, M. J. Franklin, A. Ghodsi, M. Zaharia, Spark sql: Relational data processing in spark, in: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, SIGMOD '15, ACM, New York, NY, USA, 2015, pp. 1383–1394.

[15] R. S. Xin, J. E. Gonzalez, M. J. Franklin, I. Stoica, Graphx: A resilient distributed graph system on spark, in: First International Workshop on Graph Data Management Experiences and Systems, GRADES '13, ACM, New York, NY, USA, 2013, pp. 2:1–2:6.

[16] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, G. Czajkowski, Pregel: a system for large-scale graph processing, in: Proceedings of the 2010 ACM SIGMOD International Conference on Management of data, ACM, 2010, pp. 135–146.

[17] A. Toshniwal, S. Taneja, A. Shukla, K. Ramasamy, J. M. Patel, S. Kulkarni, J. Jackson, K. Gade, M. Fu, J. Donham, N. Bhagat, S. Mittal, D. Ryaboy, Storm@twitter, in: Proceedings of the 2014 ACM SIGMOD International Conference on Management of

Data, SIGMOD '14, ACM, New York, NY, USA, 2014, pp. 147–156.

[18] P. Hunt, M. Konar, F. P. Junqueira, B. Reed, Zookeeper: Wait-free coordination for internet-scale systems, in: Proceedings of the 2010 USENIX Conference on USENIX Annual Technical Conference, USENIXATC'10, USENIX Association, Berkeley, CA, USA, 2010, pp. 11–11.

[19] A. Alexandrov, R. Bergmann, S. Ewen, J.-C. Freytag, F. Hueske, A. Heise, O. Kao, M. Leich, U. Leser, V. Markl, F. Naumann, M. Peters, A. Rheinländer, M. J. Sax, S. Schelter, M. Höger, K. Tzoumas, D. Warneke, The stratosphere platform for big data analytics, The VLDB Journal 23 (6) (2014) 939–964.

[20] Y. Yao, J. Wang, B. Sheng, J. Lin, N. Mi, Haste: Hadoop yarn scheduling based on task-dependency and resource-demand, in: Proceedings of the 2014 IEEE International Conference on Cloud Computing, CLOUD '14, IEEE Computer Society, Washington, DC, USA, 2014, pp. 184–191. doi:10.1109/CLOUD.2014.34.

[21] J. Lin, C. Dyer, Data-Intensive Text Processing with MapReduce, Morgan and Claypool Publishers, 2010.

[22] F. Bajaber, R. Elshawi, O. Batarfi, A. Altalhi, A. Barnawi, S. Sakr, Big data 2.0 processing systems: Taxonomy and open challenges, Journal of Grid Computing 14 (3) (2016) 379–405.

[23] Y. Gupta, Kibana Essentials, Packt Publishing, 2015.

[24] S. Wadkar, M. Siddalingaiah, Apache Ambari, Apress, Berkeley, CA, 2014, pp. 399–401.

[25] D. Eadline, Hadoop 2 Quick-Start Guide: Learn the Essentials of Big Data Computing in the Apache Hadoop 2 Ecosystem, 1st Edition, Addison-Wesley Professional, 2015.

[26] S. Landset, T. M. Khoshgoftaar, A. N. Richter, T. Hasanin, A survey of open source tools for machine learning with big data in the hadoop ecosystem, Journal of Big Data 2 (1) (2015) 1–36. doi:10.1186/s40537-015-0032-1.

[27] S. Chakrabarti, E. Cox, E. Frank, R. H. Gting, J. Han, X. Jiang, M. Kamber, S. S. Lightstone, T. P. Nadeau, R. E. Neapolitan, D. Pyle, M. Refaat, M. Schneider, T. J. Teorey, I. H. Witten, Data Mining: Know It All, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008.

[28] E. R. Sparks, A. Talwalkar, V. Smith, J. Kottalam, X. Pan, J. E. Gonzalez, M. J. Franklin, M. I. Jordan, T. Kraska, MLI: an API for distributed machine learning, in: 2013 IEEE 13th International Conference on Data Mining, Dallas, TX, USA, December 7-10, 2013, 2013, pp. 1187–1192.

[29] C. Giatsidis, D. M. Thilikos, M. Vazirgiannis, Evaluating cooperation in communities with the k-core structure, in: Proceedings of the 2011 International Conference on Advances in Social Networks Analysis and Mining, ASONAM '11, IEEE Computer Society, Washington, DC, USA, 2011, pp. 87–93.

[30] J. I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, A. Vespignani, K-core decomposition of internet graphs: hierarchies, self-similarity and measurement biases, NHM 3 (2) (2008) 371–393.

[31] C. Huttenhower, S. O. Mehmood, O. G. Troyanskaya, Graphle: Interactive exploration of large, dense graphs, BMC Bioinformatics 10 (2009) 417.

[32] B. Elser, A. Montresor, An evaluation study of bigdata frameworks for graph processing, in: IEEE International Conference on Big Data, 2013, pp. 60–67.

[33] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, K. Tzoumas, Apache flink^{TM}: Stream and batch processing in a single engine, IEEE Data Eng. Bull. 38 (4) (2015) 28–38.

[34] S. Aridhi, E. M. Nguifo, Big graph mining: Frameworks and techniques, Big Data Research (2016) (in press).

[35] Y. Low, D. Bickson, J. Gonzalez, et al., Distributed graphlab: A framework for machine learning and data mining in the cloud, Proc. VLDB Endow. 5 (8) (2012) 716–727.

[36] S. Aridhi, A. Montresor, Y. Velegrakis, Bladyg: A novel block-centric framework for the analysis of large dynamic graphs, in: Proceedings of the ACM Workshop on High Performance Graph Processing, HPGP '16, ACM, New York, NY, USA, 2016, pp. 39–42.

[37] B. Shao, H. Wang, Y. Li, Trinity: A distributed graph engine on a memory cloud, in: Proc. of the Int. Conf. on Management of Data, ACM, 2013.

[38] F. Zhang, J. Cao, S. U. Khan, K. Li, K. Hwang, A task-level adaptive mapreduce framework for real-time streaming data in healthcare applications, Future Gener. Comput. Syst. 43 (C) (2015) 149–160.

[39] Y. Bu, B. Howe, M. Balazinska, M. D. Ernst, The haloop approach to large-scale iterative data analysis, The VLDB Journal 21 (2) (2012) 169–190.

[40] J. Ekanayake, H. Li, B. Zhang, T. Gunarathne, S.-H. Bae, J. Qiu, G. Fox, Twister: a runtime for iterative mapreduce, in: Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing, ACM, 2010, pp. 810–818.

[41] P. Resnick, H. R. Varian, Recommender systems, Commun. ACM 40 (3) (1997) 56–58.

[42] J. Domann, J. Meiners, L. Helmers, A. Lommatzsch, Real-time news recommendations using apache spark, in: Working Notes of CLEF 2016 - Conference and Labs of the Evaluation forum, Évora, Portugal, 5-8 September, 2016., 2016, pp. 628–641.

[43] G. Bello-Orgaz, J. J. Jung, D. Camacho, Social big data: Recent achievements and new challenges, Information Fusion 28 (2016) 45 – 59.

[44] C. Yin, Z. Xiong, H. Chen, J. Wang, D. Cooper, B. David, A literature survey on smart cities, Science China Information Sciences 58 (10) (2015) 1–18.

[45] C. L. Stimmel, Building Smart Cities: Analytics, ICT, and Design Thinking, Auerbach Publications, Boston, MA, USA, 2015.

[46] G. Piro, I. Cianci, L. Grieco, G. Boggia, P. Camarda, Information centric services in smart cities, Journal of Systems and Software 88 (2014) 169 – 188.

[47] X. Xu, Q. Z. Sheng, L. J. Zhang, Y. Fan, S. Dustdar, From big data to big service, Computer 48 (7) (2015) 80–83.